arXiv:1404.0062v1 [cs.IT] 31 Mar 2014

On redundancy of memoryless sources over countable alphabets

Maryam Hosseini Narayana Santhanam

April 2, 2014

Abstract

The minimum average number of bits need to describe a random variable is its entropy. This supposes knowledge of the distribution of the random variable. On the other hand, *universal compression* supposes that the distribution of the random variable, while unknown, belongs to a known set $\mathcal{P}$ of distributions. Such *universal descriptions* for the random variable are agnostic to the identity of the distribution in $\mathcal{P}$. But because they are not matched exactly to the underlying distribution of the random variable, the average number of bits they use is higher, and the excess over the entropy used is the *redundancy*. This formulation is fundamental to problems not just in compression, but also estimation and prediction and has a wide variety of applications from language modeling to insurance.

In this paper, we study the redundancy of universal encodings of strings generated by independent identically distributed (*i.i.d.*) sampling from a set $\mathcal{P}$ of distributions over a countable support. We first show that if describing a single sample from $\mathcal{P}$ incurs finite redundancy, then $\mathcal{P}$ is tight but that the converse does not always hold.

If a single sample can be described with finite *worst-case-regret* (a more stringent formulation than redundancy above), then it is known that describing length-n *i.i.d.* samples only incurs a diminishing (in n) redundancy *per symbol* as n increases. However, we show it is possible that a collection $\mathcal{P}$ incurs finite redundancy, yet description of length-n *i.i.d.* samples incurs a constant redundancy per symbol encoded. We then show a sufficient condition on $\mathcal{P}$ such that length-n *i.i.d.* samples will incur diminishing redundancy per symbol encoded.

A number of statistical inference problems of significant contemporary interest, such as text classification, language modeling, and DNA microarray analysis, requires computing inferences based on observed sequences of symbols in which the sequence length or sample size is comparable or even smaller than the set of symbols, the *alphabet*. For instance, language models for speech recognition estimate distributions over English words using text examples much smaller than the vocabulary.

To model these problems, several lines of work have considered universal compression over large alphabets. Generally, the idea here is to model the problem at hand with a class of models $\mathcal{P}$ instead of a single distribution. The model underlying the data is assumed or known to belong to the class $\mathcal{P}$, but the exact identity of the model remains unknown. Instead, we aim to use a universal description of data.

The universal description uses more bits on an average (over the sample) than if the underlying model was known, and the additional number of bits used by the universal description is called the *redundancy* against the true model. The average excess bits over the entropy of the model will be refered to as the *model redundancy* for that model. Since one does not know the true model in general, a common approach is to consider *strong redundancy*, which is the supremum over all models of the class of the model redundancy.

Typically we look at sequences of *i.i.d.* symbols, and therefore we usually refer to the redundancy of distributions over length-n sequences obtained by *i.i.d.* sampling from distributions from $\mathcal{P}$. The length n of sequences considered will typically be refered to as the sample size.

The nuances of prediction, compression or estimation where the alphabet size and sample size are roughly equal are not well captured by the asymptotics of the strong redundancy with of a class over a fixed alphabet and sample size going to infinity. Rather, they are better captured when we begin with a countably infinite support and let the sample size approach infinity, or when we let the alphabet size scale as a function of the sample size.

To begin with, the collection of all *i.i.d.* distributions over countably infinite supports or alll *i.i.d.* distributions over an alphabet whose size is comparable to the sample length of interest have very high redundancy that renders most estimation or prediction problems impossible. Therefore, there are several alternative formulations to tackle language modeling, or classification and estimation questions over large alphabets.

Patterns One line of work is the patterns [1] approach that considers the compression of the *pattern* of a sequence rather than the sequence itself. Patterns abstract the identities of symbols, and indicate only the relative order of appearance. For example, the pattern of TATTLE is 121134, while that of HONOLULU is 12324545. The point to note is that patterns of length-n *i.i.d.* sequences can be compressed (no matter what the underlying countably infinite alphabet is) with redundancy that grows sublinearly in n, therefore the excess bits needed is asymptotically diminishing (in n) per-symbol redundancy. Indeed insights learnt in this line of work will be used to understand compression of sequences as well in this paper.

Envelope on model classes A second line of work considers restricted model classes for applications, particularly where the collection of models can be described in terms of an envelope [2]. This approach leads to an understanding of the worst case formulations. In particular, we are interested in the result that if the *worst-case* redundancy (different from, and a more stringent formulation than the strong redundancy described here) of describing a single sample is finite, then the per-symbol strong redundancy diminishes to 0. We will interpret this result towards the end of the introduction.

Data derived consistency A third line of work ignores the uniform convergence framework underlying strong redundancy formulations. This is useful for large or infinite alphabet model classes which have poor or no strong redundancy guarantees, but ask a question that cannot be answered with the patterns approach above. In this line of work, one obtains results on the *model redundancy* described above instead of *strong* redundancy . For example, a model class is said to be weakly compressible if there is a universal measure that ensures that for all models, the model redundancy normalized by the sample size (per-symbol) diminishes to 0. The rate at which the per-symbol model redundancy diminishes to 0 depends on the underlying model, and for some models could be arbitrarily slower than others. Given a particular blocklength, n, there may be hence no non-trivial guarantee that holds over the entire model class unlike the strong redundancy formulation.

But if we add on the additional constraint that we should be estimate the rate of convergence from the data, we get the *data-derived* consistency formulations in [3]. Fundamental to further research in this direction is a better understanding of how single letter redundancy (of $\mathcal{P}$) relates to

the redundancy of length-n strings (that of $\mathcal{P}^n$). The primary theme of this paper is to collect such results on strong redundancy of classes over countably infinite support.

In the fixed alphabet setting, this connection is well understood. If the alphabet has size k, the redundancy of $\mathcal{P}$ is easily seen to be always finite (in fact $\leq \log k$) and that of $\mathcal{P}^n$ scales as $\frac{k}{2} \log n$. But when $\mathcal{P}$ does not have a finite support, the above bounds are meaningless.

On the other hand, the redundancy of a class $\mathcal{P}$ over $\mathbb{N}$ may in general be infinite. But what about the case where redundancy of $\mathcal{P}$ is finite? Now a well known *redundancy-capacity* [4] argument can be used to interpret the redundancy—and this tells us that the redundancy is the amount of information we can get about the source from the data. In this case, finite (infinite respectively) redundancy of $\mathcal{P}$ implies that a single symbol contains finite (infinite respectively) amount of information about the model.

The natural question then is does it imply that the redundancy of length-n *i.i.d.* strings from $\mathcal{P}$ grows sublinearly? Equivalently, do finite redundancy classes over $\mathbb{N}$ behave essentially like some of their fixed alphabet counterparts? In some formulations (the worst-case), this indeed holds. Roughly speaking, this informs us that as the universal encoder sees more and more of the sequence, it learns less and less of the underlying model. This would be in line with our intuition where seeing more data fixes the model, so as we so more data there is less to learn.

To understand these connections, we first show that if the redundancy of a collection $\mathcal{P}$ of distributions over $\mathbb{N}$ is finite, then $\mathcal{P}$ is tight. This turns out to be a useful tool to check if the redundancy is finite in [3] for example.

But in a departure from previous formulations, we then demonstrate that it is possible for a class $\mathcal{P}$ to have finite redundancy, yet the redundancy of length-n strings sampled *i.i.d.* from $\mathcal{P}$ does not grow sublinearly in n! Therefore, roughly speaking, no matter how much of the sequence the universal encoder has seen, it learns at least a constant number of bits about the underlying model each time it sees another symbol. No matter how much data we see, there is more to learn! We finally obtain a sufficient condition on a class $\mathcal{P}$ such that the asymptotic per-symbol redundancy of length-n *i.i.d.* strings diminishes to 0.

1 Notation and background

We will review the notions of universal compression, redundancy and patterns here, as well as some allied results that we will make use of in this paper.

Let $\mathcal{P}$ be a collection of distributions over $\mathbb{N}$. Let $\mathcal{P}^n$ be the set of distributions over length-n sequences obtained by *i.i.d.* sampling from $\mathcal{P}$. Let $\mathcal{P}^\infty$ be the collection of measures over infinite length sequences[1] of $\mathbb{N}$ obtained by *i.i.d.* sampling from distributions of $\mathcal{P}$.

[1] Observe that $\mathcal{N}^n$ is countable for every n. For simplicity of exposition, we will think of each length-n string $\mathbf{x}$ as a subset of $\mathcal{N}^\infty$—the set of all semi-infinite strings of naturals that begin with $\mathbf{x}$. Each subset of $\mathcal{N}^n$ is therefore a subset of $\mathcal{N}^\infty$.

Now the collection $\mathcal{J}$ of all subsets of $\mathcal{N}^n$, $n \geq 1$, is a semi-algebra [5]. The probabilities *i.i.d.* sampling assigns to finite unions of disjoint sets in $\mathcal{J}$ is the sum of that assigned to the components of the union. Therefore, there is a sigma-algebra over the uncountable set $\mathcal{N}^\infty$ that extends $\mathcal{J}$ and matches the probabilities assigned to sets in $\mathcal{J}$ by *i.i.d.* sampling. The reader can assume the sigma-algebra is the minimal such extension. $\mathcal{P}^\infty$ is the measure on this sigma-algebra that matches what the probabilities *i.i.d.* sampling gives to sets in $\mathcal{J}$. See, *e.g.* [5], for a development of elementary measure theory that lays out the above results.

1.1 Redundancy

1.1.1 Strong compression

A class $\mathcal{P}^\infty$ of measures over infinite sequences of natural numbers is called strongly compressible if there is a measure q over $\mathbb{N}^\infty$ satisfying

$$\limsup_{n\to\infty} \sup_{p\in\mathcal{P}^\infty} \frac{1}{n} E_p \log \frac{p(X^n)}{q(X^n)} = 0. \tag{1}$$

In particular, we call

$$\inf_q \sup_{p\in\mathcal{P}^\infty} \frac{1}{n} E_p \log \frac{p(X^n)}{q(X^n)}$$

the redundancy of length-n sequences, or length-n redundancy. The single letter redundancy refers to the special case when $n = 1$.

We will often be concerned with sequences of symbols drawn *i.i.d.* from $\mathcal{P}$, and let $\mathcal{P}^\infty$ be the measures induced on infinite sequences of naturals obtained by *i.i.d.* sampling from distributions in $\mathcal{P}$. Our primary goal is to understand the connections between the single letter redundancy on the one hand and the behavior of length-n *i.i.d.* redundancy on the other. Strong compression length-n redundancy can be seen as the capacity of a channel from $\mathcal{P}$ to $\mathbb{N}^n$, where the conditional probability distribution over $\mathbb{N}^n$ given $p \in \mathcal{P}$ is simply the distribution p over length-n sequences.

We note that it is possible to define an even more stringent notion—a *worst case* formulation. For length-n sequences, this is

$$\inf_q \sup_{p\in\mathcal{P}^\infty} \frac{1}{n} \sup \log \frac{p(X^n)}{q(X^n)}.$$

We will not concern ourselves with the worst case formulation in this paper, but mention it in passing for comparisons. In the worst case setting, finite single letter redundancy is necessary and sufficient for the asymptotic per-symbol worst case redundancy to diminish to 0.

But we show in this paper that it is not necessarily the case for strong redundancy. It is quite possible that classes with finite single letter strong redundancy have asymptotic per-symbol strong redundancy bounded away from 0.

1.2 Patterns

Recent work [1] has formalized a similar framework for countably infinite alphabets. This framework is based on the notion of *patterns* of sequences that abstract the identities of symbols, and indicate only the relative order of appearance. For example, the pattern of PATTERN is 1233456. The k'th distinct symbol of a string is given an index k when it first appears, and that index is used every time the symbol appears henceforth. The crux of the patterns approach is to consider the set of measures induced over patterns of the sequences instead of considering the set of measures $\mathcal{P}$ over infinite sequences,

Denote the pattern of a string $\mathbf{x}$ by $\Psi(\mathbf{x})$. There is only one possible pattern of strings of length 1 (no matter what the alphabet, the pattern of a length-1 string is 1), two possible patterns of strings of length 2 (11 and 12), and so on. The number of possible patterns of length n is the n'th Bell number [1] and we denote the set of all possible length n patterns by Ψ^n. The measures induced

on patterns by a corresponding measure p on infinite sequences of natural numbers assigns to any pattern ψ a probability

$$p(\psi) = p(\{\mathbf{x} : \Psi(\mathbf{x}) = \psi\}).$$

In [1] the length-n pattern redundancy,

$$\inf_q \sup_{p \in \mathcal{P}^\infty} \frac{1}{n} E_p \log \frac{p(\Psi(X^n))}{q(\Psi(X^n))},$$

was shown to be upper bounded by $\pi(\log e)\sqrt{\frac{2n}{3}}$. It was also shown in [6] that there is a measure q over infinite length sequences which satisfies for all n simultaneously

$$\sup_{p \in \mathcal{P}^\infty} \frac{1}{n} E_p \log \frac{p(\Psi(X^n))}{q(\Psi(X^n))} \leq \pi(\log e)\sqrt{\frac{2n}{3}} + \log(n(n+1)).$$

Let the measure induced on patterns by q be denoted as q_Ψ for convenience.

We can (naively) interpret the probability estimator q_Ψ as a sequential prediction procedure that estimates the probability that the symbol X_{n+1} will be "new" (has not appeared in X_1^n), and the probability that X_{n+1} takes a value that has been seen so far. This view of estimation also appears in the statistical literature on Bayesian nonparametrics that focuses on exchangeability. Kingman [7] advocated the use of *exchangeable random partitions* to accommodate the analysis of data from an alphabet that is not bounded or known in advance. A more detailed discussion of the history and philosophy of this problem can be found in the works of Zabell [8, 9] collected in [10].

1.3 Cummulative distributions and tight classes

For our purposes, the cumulative distribution function of any probability distribution p on $\mathbb{N}$ is a function $F_p : \mathbb{R}^+ \cup \{\infty\} \to [0,1]$ defined in the following (slightly unconventional) way. We obtain F_p by first defining F_p on points in the support of p in the way cumulative distribution functions are normally defined. We define F_p for all other nonnegative real numbers by linearly interpolating between the values in the support of p. Finally, $F_p(\infty) := 1$.

Let $F_p^{-1} : [0,1] \mapsto \mathbb{R}^+ \cup \{\infty\}$ denote the inverse function of F_p. Then $F_p^{-1}(x) = 0$ for all $0 \leq x < F_p(0)$. If p has infinite support then $F_p^{-1}(1) = \infty$, else $F_p^{-1}(1)$ is the smallest natural number y such that $F_p(y) = 1$.

A class $\mathcal{P}$ of distributions on $\mathbb{N}$ is defined to be *tight* if for all $\gamma > 0$,

$$\sup_{p \in \mathcal{P}} F_p^{-1}(1-\gamma) < \infty.$$

2 Redundancy and tightness

We focus on the single letter redundancy in this section, and explore the connections between the single letter redundancy of a class $\mathcal{P}$ and the tightness of $\mathcal{P}$.

Lemma 1. A class $\mathcal{P}$ with bounded strong redundancy is tight. Namely, if the strong redundancy of $\mathcal{P}$ is finite, then for any $\gamma > 0$

$$\sup_{p \in \mathcal{P}} F_p^{-1}(1-\gamma) < \infty.$$

Proof $\mathcal{P}$ has bounded strong redundancy. Let q be a distribution over $\mathbb{N}$ such that

$$\sup_{p\in\mathcal{P}} D(p||q) < \infty,$$

and we define $R = \sup_{p\in\mathcal{P}} D(p||q)$. It follows that for all $p \in \mathcal{P}$ and any m,

$$p(\left|\log \frac{p(X)}{q(X)}\right| > m) \leq (R + (2\log e)/e)/m,$$

To see the above, note that if S is the set of all numbers such that $p(x) < q(x)$, a well-known convexity argument shows that

$$\sum_x p(x) \log \frac{p(x)}{q(x)} \geq p(S) \log \frac{p(S)}{q(S)} \geq -\frac{\log e}{e}.$$

We prove the lemma by contradiction. Pick m so large that $(R + (2\log e)/e)/m < \gamma/2$. For all p, we show that

$$p\big(x : x \geq F_q^{-1}(1 - \gamma/2^{m+1})\big) \leq \gamma.$$

To see the above, observe that we can split the tail $x \geq F_q^{-1}(1-\gamma/2^{m+1})$ into two parts—(i) numbers x such that $\log \frac{p(x)}{q(x)} > m$. This set has probability $< \gamma/2$ under p. (ii) remaining numbers x such that $\log \frac{p(x)}{q(x)} < m$. This set has probability $\leq \gamma/2^{m+1}$ under q, and therefore probability $\leq \gamma/2$ under p. The lemma follows. □

The converse is not necessarily true. Tight classes need not have finite single letter redundancy as the following example demonstrates.

Construction Consider the following class $\mathcal{I}$ of distributions over $\mathbb{N}$. First partition the set of natural numbers into the sets T_i, $i \geq 0$, where

$$T_i = \{2^k, \ldots, 2^{k+1} - 1\}.$$

Note that $|T_k| = 2^k$. Now, $\mathcal{I}$ is the collection of all possible distributions that can be formed as follows. For all $i \geq 1$, we pick exactly one element of T_i and assign it probability $1/(i(i+1))$. Note that the set $\mathcal{I}$ is uncountably infinite. □

Corollary 2. The set $\mathcal{I}$ of distributions is tight.
Proof For all $p \in \mathcal{I}$,

$$\sum_{n\geq 2^k} p(n) = \frac{1}{k+1},$$

namely, all tails are uniformly bounded over the class $\mathcal{I}$ to ensure insurability. Put another way, for all $\delta > 0$ and all distributions $p \in \mathcal{I}$,

$$F_p^{-1}(1-\delta) \leq 2^{\lfloor \frac{1}{\delta} \rfloor} - 1.$$ □

On the other hand,

Proposition 1. The collection $\mathcal{I}$ does not have finite redundancy.

Proof Suppose q is any distribution over $\mathbb{N}$. We will show that $\exists p \in \mathcal{I}$ such that

$$\sum_{n\geq 1} p(n) \log \frac{p(n)}{q(n)}$$

is not finite. Since the entropy of every $p \in \mathcal{I}$ is finite, we just have to show that for any distribution q over $\mathbb{N}$, there $\exists p \in \mathcal{I}$ such that

$$\sum_{n\geq 1} p(n) \log \frac{1}{q(n)}$$

is not finite.

Consider any distribution q over $\mathbb{N}$. Observe that for all i, $|T_i| = 2^i$. It follows that for all i there is $x_i \in T_i$ such that

$$q(x_i) \leq \frac{1}{2^i}.$$

But by construction, $\mathcal{I}$ contains a distribution p that has for its support $\{x_i : i \geq 1\}$ identified above. Furthermore p assigns

$$p(x_i) = \frac{1}{i(i+1)} \qquad \forall\, i \geq 1.$$

The KL divergence from p to q is not finite and the Lemma follows. □

3 Length-n redundancy

We study how the single letter properties of a collection $\mathcal{P}$ of distributions influences the compression of length-n strings obtained by *i.i.d.* sampling from distributions in $\mathcal{P}$. Namely, we try to characterize when the length-n redundancy of $\mathcal{P}^\infty$ grows sublinearly in the blocklength n.

Lemma 3. Let $\mathcal{P}$ be a class of distributions over a countable support $\mathcal{X}$. For some $m \geq 1$, consider m pairwise disjoint subsets $S_i \subset \mathcal{X}$ $(1 \leq i \leq m)$ and let $\delta > 1/2$. If there exist $p_1, \ldots, p_m \in \mathcal{P}$ such that

$$p_i(S_i) \geq \delta,$$

then for all distributions q over $\mathcal{X}$,

$$\sup_{p\in\mathcal{P}} D(p||q) \geq \delta \log m.$$

In particular if there are an infinite number of sets S_i, $i \geq 1$ and distributions $p_i \in \mathcal{P}$ such that $p_i(S_i) \geq \delta$, then the redundancy is infinite.

Proof This is a simplified formulation of the *distinguishability* concept in [4]. For a proof, see *e.g.* [11]. □

3.1 Counterexample

We now show that it is possible for the single letter redundancy of a class $\mathcal{B}$ of distributiosn to be finite, yet the asymptotic per-symbol redundancy (the length-n redundancy of $\mathcal{B}^\infty$ normalized by n) remains bounded away from 0 in the limit the blocklength goes to infinity. To show this, we obtain such a class $\mathcal{B}$.

Construction As before partition the set $\mathbb{N}$ into $T_i = \{2^i, \ldots, 2^{i+1} - 1\}$, $i \geq 0$. Recall that T_i has 2^i elements. For all $\epsilon > 0$, let $n_\epsilon = \lfloor \frac{1}{\epsilon} \rfloor$. Let $1 \leq j \leq 2^{n_\epsilon}$ and let $p_{\epsilon,j}$ be a distribution on $\mathbb{N}$ that assigns probability $1 - \epsilon$ to the number 1 (or equivalently, to the set T_0), and ϵ to the j'th smallest element of T_{n_ϵ}, namely the number $2^{n_\epsilon} + j - 1$. $\mathcal{B}$ (mnemonic for binary, since every distribution has at support of size 2) is the collection of distributions $p_{\epsilon,j}$ for all $\epsilon > 0$ and $1 \leq j \leq 2^{n_\epsilon}$. $\mathcal{B}^\infty$ is the set of measures over infinite sequences of numbers corresponding to *i.i.d.* sampling from $\mathcal{B}$. □

Proposition 2. For all $n \geq 1$,

$$\inf_q \sup_{p \in B^\infty} E_p \log \frac{p(X^n)}{q(X^n)} \geq n\left(1 - \frac{1}{n}\right)^n.$$

Proof For all n, define 2^n pairwise disjoint sets S_i of $\mathbb{N}^n$, $1 \leq i \leq 2^n$, where

$$S_i = \{1, 2^n + i - 1\}^n - \{1^n\}$$

is the set of all length-n strings containing at most two numbers (1 and $2^n + i - 1$) and at least one occurance of $2^n + i - 1$. Clearly, for distinct i and j between 1 and 2^n. S_i and S_j are disjoint. Furthermore, the distribution $p_{\frac{1}{n},i}$ assigns S_i the probability

$$p_{\frac{1}{n},i}(S_i) = 1 - \left(1 - \frac{1}{n}\right)^n > 1 - \frac{1}{e}.$$

From Lemma 3, it follows that length-n redundancy of $\mathcal{B}^\infty$ is lower bounded by

$$\left(1 - \frac{1}{e}\right) \log 2^n = n\left(1 - \frac{1}{e}\right).$$ □

3.2 Sufficient condition

In this section, we show a sufficient condition on single letter marginals of $\mathcal{P}$ and its redundancy that allows for *i.i.d.* length-n redundancy of $\mathcal{P}^\infty$ to grow sublinearly with n. This condition is, however, not necessary—and the characterization of a condition that is both necessary and sufficient is as yet open.

For all $\epsilon > 0$, let $A_{p,\epsilon}$ is the set of all elements in the support of p which have probability $\geq \epsilon$, and let $T_{p,\epsilon} = \mathbb{N} - A_{p,\epsilon}$ For all i, the sets

$$G_i = \{x^i : A_{p,\frac{2\log i}{i}} \subseteq \{x_1, x_2, \ldots, x_i\}\}$$

where in a minor abuse of notation, we use $\{x_1, \ldots, x_i\}$ to denote the set of distinct symbols in the string x_1^i. Let $B_i = \mathbb{N} - G_i$. Observe from an argument similar to the coupon collector problem [] that (correct for log 2 bases)

Lemma 4. For all $i \geq 1$,

$$p(B_i) \leq \frac{j}{2\log j}\left(1 - \frac{2\log j}{j}\right)^j \leq \frac{1}{j \log j}.$$ □

Theorem 5. Suppose $\mathcal{P}$ is a class of distributions over $\mathbb{N}$. Let the entropy of $p \in \mathcal{P}$, denoted by $H(p)$, be uniformly bounded over the entire class, and in addition let the redundancy of the class be finite. Namely,

$$\sup_{p\in\mathcal{P}} \sum_{x\in\mathbb{N}} p(x) \log \frac{1}{p(x)} < \infty \quad \text{and } \exists q_1 \text{ over } \mathbb{N} \text{ s.t.} \quad \limsup_{\delta\to 0} \sup_{p\in\mathcal{P}} \sum_{x\in\mathbb{N}} p(x) \log \frac{p(x)}{q_1(x)} < \infty.$$

Recall that for any distribution p, the set $T_{p,\delta}$ denotes the support of p all of whose probabilities are $< \delta$. Let

$$\lim_{\delta\to 0} \sup_{p\in\mathcal{P}} \sum_{x\in T_{p,\delta}} p(x) \log \frac{1}{p(x)} = 0 \quad \text{and } \exists q_1 \text{ over } \mathbb{N} \text{ s.t.} \quad \limsup_{\delta\to 0} \sup_{p\in\mathcal{P}} \sum_{x\in T_{p,\delta}} p(x) \log \frac{p(x)}{q_1(x)} = 0. \tag{2}$$

Then, the redundancy of length-n distributions obtained by *i.i.d.* sampling from distributions in $\mathcal{P}$, denoted by $R_n(\mathcal{P}^\infty)$, grows sublinearly

$$\limsup_{n\to\infty} \frac{1}{n} R_n(\mathcal{P}^\infty) = 0.$$

Proof Let q_Ψ be the optimal universal pattern encoder over patterns of *i.i.d.* sequences from Section 1.2. Since the redundancy of $\mathcal{P}$ is finite, let q_1 be a universal distribution over $\mathbb{N}$ that attains finite redundancy for $\mathcal{P}$. We consider a universal encoder as follows:

$$\begin{aligned}
q(x^n) &= q(x^n, \Psi(x^n)) \\
&\overset{(a)}{=} q(\psi_1, x_1, \psi_2, x_2, \ldots, \psi_n, x_n) \\
&= \prod_{i\geq 1} q(\psi_i | \psi_1^{i-1}, x_1^{i-1}) \prod_{j\geq 1} q(x_j | \psi_1^j, x_1^{j-1}) \\
&\overset{\text{def}}{=} \prod_{i\geq 1} q_\Psi(\psi_i | \psi_1^{i-1}) \prod_{j\geq 1} q(x_j | \psi_1^j, x_1^{j-1})
\end{aligned}$$

where in (a), we denote $\Psi(x^n) = \psi_1, \ldots, \psi_n$. Furthermore we define for all $x_1^{i-1} \in \mathbb{N}^{i-1}$ and all $\psi^i \in \Psi^i$ such that $\psi^{i-1} = \Psi(x^{i-1})$,

$$q(x_i | \psi_1^i, x_1^{i-1}) \overset{\text{def}}{=} \begin{cases} 1 & \text{if } x_i \in \{x_1, \ldots, x_{i-1}\} \text{ and } \Psi(x^i) = \psi^i \\ q_1(x_i) & \text{if } x_i \notin \{x_1, \ldots, x_{i-1}\} \text{ and } \Psi(x^i) = \psi^i. \end{cases}$$

Namely, we use an optimal universal pattern encoder over patterns of *i.i.d.* sequences, and encode any new symbol using a universal distribution over $\mathcal{P}$. We now bound the redundancy of q as defined above. We have for all $p \in \mathcal{P}^\infty$,

$$\begin{aligned}
E_p \log \frac{p(X^n)}{q(X^n)} &= \sum_{x^n} p(x^n) \log \prod_{i\geq 1} \frac{p(\psi_i | \psi_1^{i-1}, x_1^{i-1})}{q_\Psi(\psi_i | \psi_1^{i-1})} \prod_{j\geq 1} \frac{p(x_j | \psi_1^j, x_1^{j-1})}{q(x_j | \psi_1^j, x_1^{j-1})} \\
&= \sum_{x^n} p(x^n) \sum_{i=1}^n \log \frac{p(\psi_i | \psi^{i-1}, x_1^{i-1})}{q_\Psi(\psi_i | \psi^{i-1})} + \sum_{x^n} p(x^n) \sum_{j=1}^n \log \frac{p(x_j | \psi_1^j, x_1^{j-1})}{q(x_j | \psi_1^j, x_1^{j-1})}
\end{aligned}$$

The first term, normalized by n, can be upper bounded by as follows

$$\frac{1}{n}\sum_{x^n} p(x^n)\sum_{i=1}^{n}\log\frac{p(\psi_i|\psi^{i-1},x_1^{i-1})}{q_\Psi(\psi_i|\psi^{i-1})} \le \frac{1}{n}\sum_{i=1}^{n}\sum_{x^i} p(x^i)\log\frac{p(\psi_i|\psi^{i-1},x_1^{i-1})}{p(\psi_i|\psi^{i-1})} + \pi\sqrt{\frac{2}{3n}}$$
$$= \frac{1}{n}\sum_{i=1}^{n}(H(\Psi_i|\Psi^{i-1}) - H(\Psi_i|X^{i-1})) + \pi\sqrt{\frac{2}{3n}}.$$

Now

$$\begin{aligned} H(\Psi_i|\Psi^{i-1}) - H(\Psi_i|X^{i-1}) &\le H(X_i|\Psi^{i-1}) - H(\Psi_i|X^{i-1}) \\ &= H - H(\Psi_i|X^{i-1}) \\ &= \sum_{x^{i-1}} p(x) \sum_{x\notin\{x_1,\ldots,x^{i-1}\}} p(x)\log\frac{1}{p(x)} \\ &\le p(G_{i-1}) \sum_{x\in T_{p,2\frac{log(i-1)}{i-1}}} p(x)\log\frac{1}{p(x)} + p(B_{i-1})H \\ &\le \sum_{x\in T_{p,2\frac{log(i-1)}{i-1}}} p(x)\log\frac{1}{p(x)} + \frac{H}{(i-1)\log(i-1)}. \end{aligned}$$

We have split the length j sequences into the sets G_j and B_j and use separate bounds on each set that hold uniformly over the entire model class. The last inequality (a) above follows from Lemma 4. From condition (2) of the Theorem, we have that

$$\lim_{i\to\infty}\sup_{p\in\mathcal{P}} \sum_{x\in T_{p,2\frac{log(i-1)}{i-1}}} p(x)\log\frac{1}{p(x)} = 0$$

$$\text{implying that } \lim_{n\to\infty}\sup_{p\in\mathcal{P}} \frac{1}{n}\sum_{i=1}^{n} \sum_{x\in T_{p,2\frac{log(i-1)}{i-1}}} p(x)\log\frac{1}{p(x)} + \frac{H}{(i-1)\log(i-1)} = 0.$$

For the second term, we have

$$\begin{aligned} \sum_{x^n} p(x^n)\sum_{j=1}^{n}\log\frac{p(x_j|\psi_1^j,x_1^{j-1})}{q(x_j|\psi_1^j,x_1^{j-1})} &= \sum_{j=1}^{n}\sum_{x^j} p(x^j)\log\frac{p(x_j|\psi_1^j,x_1^{j-1})}{q(x_j|\psi_1^j,x_1^{j-1})} \\ &\le \sum_{j=1}^{n}\sum_{x^{j-1}} p(x^{j-1}) \sum_{x_j\notin\mathcal{A}(x^{j-1})} p(x_j)\log\frac{p(x_j)}{q_1(x_j)} \\ &\le \sum_{j=1}^{n}\left(p(G_j) \sum_{x_j\notin A_{p,\frac{2\log j}{j}}} p(x_j)\log\frac{p(x_j)}{q_1(x_j)} + Rp(B_j)\right) \\ &\le \sum_{j=1}^{n}\left(\sum_{x_j\notin A_{p,\frac{2\log j}{j}}} p(x_j)\log\frac{p(x_j)}{q_1(x_j)} + \frac{R}{j\log j}\right). \end{aligned}$$

where as before, the last inequality is from Lemma 4. Again from condition (2), we have

$$\sup_{p\in\mathcal{P}} \sum_{x_j\notin A_{p,\frac{2\log j}{j}}} p(x_j)\log\frac{p(x_j)}{q_1(x_j)} + \frac{R}{j\log j} = o(1)$$

Therefore

$$\sup_{p\in\mathcal{P}} \frac{1}{n}\sum_{j=1}^{n}\left(\sum_{x_j\notin A_{p,\frac{2\log j}{j}}} p(x_j)\log\frac{p(x_j)}{q_1(x_j)} + \frac{R}{j\log j}\right) = o(1)$$

as well. The theorem follows. □

A few comments about (2) in Theorem 5 are in order. Neither condition automatically implies the other. The set $\mathcal{B}$ of distributions in Section 3.1 is an example where every distribution has finite entropy, the redundancy of $\mathcal{B}$ is finite,

$$\lim_{\delta\to 0}\sup_{p\in\mathcal{B}}\sum_{x\in T_{p,\delta}} p(x)\log\frac{1}{p(x)} = 0 \quad \text{but } \forall q \text{ over } \mathbb{N} \text{ s.t.} \quad \lim_{\delta\to 0}\sup_{p\in\mathcal{P}}\sum_{x\in T_{p,\delta}} p(x)\log\frac{p(x)}{q_1(x)} > 0.$$

We will now construct another set $\mathcal{U}$ of distributions over $\mathbb{N}$ such that every distribution in $\mathcal{U}$ has finite entropy, the redundancy of $\mathcal{U}$ is finite,

$$\lim_{\delta\to 0}\sup_{p\in\mathcal{U}}\sum_{x\in T_{p,\delta}} p(x)\log\frac{1}{p(x)} > 0 \quad \text{but } \forall q \text{ over } \mathbb{N} \text{ s.t.} \quad \lim_{\delta\to 0}\sup_{p\in\mathcal{P}}\sum_{x\in T_{p,\delta}} p(x)\log\frac{p(x)}{q_1(x)} = 0. \tag{3}$$

At the same time, the length-n redundancy of $\mathcal{U}^\infty$ diminishes sublinearly. This is therefore also an example to show that the conditions in Theorem 5 are only sufficient, but in fact not necessary. It is yet open to find a condition on single letter marginals that is both necessary and sufficient for the asymptotic per-symbol redundancy to diminish to 0.

Construction $\mathcal{U}$ is a countable collection of distributions p_k, $k\geq 1$ where

$$p_k(x) = \begin{cases} 1-\frac{1}{k^2} & x = 0, \\ \frac{1}{k^2 2^{k^2}} & 1\leq x\leq 2^{k^2}. \end{cases}$$ □

The entropy of $p_k\in\mathcal{U}$ is therefore $1+h\big(\frac{1}{k^2}\big)$. Note that the redundancy of $\mathcal{U}$ is finite too. To see this, note that

$$\sum_{n\geq 1}\sup_{k\geq 1} p_k(n) \leq \sum_{n\geq 1}\sum_{p_k:k\geq 1} p_k(n) = \sum_{p_k:k\geq 1}\sum_{n\geq 1} p_k(n) = \sum_{p_k:k\geq 1}\frac{1}{k^2} = \frac{\pi^2}{6}. \tag{4}$$

Furthermore, letting $R\stackrel{\text{def}}{=}\log\Big(\sum_{n\geq 1}\sup_{k\geq 1} p_k(n)\Big)$, we have that the distribution

$$q(n) = \begin{cases} 1/2 & n = 0, \\ \frac{\sup_{k\geq 1} p_k(n)}{2^{R+1}} & n\geq 1. \end{cases}$$

satisfies for all $p_k \in \mathcal{U}$

$$\sum_{n \geq 0} p_k(n) \log \frac{p_k(n)}{q(n)} \leq 1 + \frac{R+1}{k^2} \leq R + 2,$$

implying that the redundancy is $\leq R + 2$. Furthermore, (4) implies from [2] that the length-n redundancy of $\mathcal{U}^\infty$ diminishes sublinearly. Now pick an integer $m \geq 1$. We have for all $p \in \mathcal{U}$,

$$\sum_{n \in T_{p, \frac{1}{m^2 2^{m^2}}}} p(n) \log \frac{p(n)}{q(n)} \leq \frac{R+1}{m^2},$$

yet for all $k \geq m$, we have

$$\sum_{n \in T_{p, \frac{1}{m^2 2^{m^2}}}} p_k(n) \log \frac{1}{p_k(n)} = 1.$$

Thus it is easy to see that $\mathcal{U}$ indeed satisfies (3).

References

[1] A. Orlitsky, N.P. Santhanam, and J. Zhang. Universal compression of memoryless sources over unknown alphabets. *IEEE Transactions on Information Theory*, 50(7):1469—1481, July 2004.

[2] S. Boucheron, A. Garivier, and E. Gassiat. Coding on countably infinite alphabets. Available from arXiv doc id: 0801.2456, 2008.

[3] N. Santhanam, V. Anantharam, A. Kavcic, and W. Szpankowski. Data driven weak universal redundancy. Submitted for publication. Full version available form arXiv.

[4] N. Merhav and M. Feder. Universal prediction. *IEEE Transactions on Information Theory*, 44(6):2124—2147, October 1998.

[5] Jeffrey Rosenthal. *A first look at rigorous probability theory*. World Scientific, 2nd edition, 2008.

[6] Narayana Santhanam. *Probability estimation and compression involving large alphabets*. PhD thesis, University of California, San Diego, 2006.

[7] J.F.C. Kingman. The mathematics of genetic diversity. *SIAM*, 1980.

[8] S.L. Zabell. Predicting the unpredictable. *Synthese*, 90:205–232, 1992.

[9] S.L. Zabell. The continuum of inductive methods revisited. In John Earman and John D. Norton, editors, *The Cosmos of Science: Essays of Exploration*, chapter 12. The University of Pittsburgh Press, Pittsburgh, PA, USA, 1997.

[10] S.L. Zabell. *Symmetry and Its Discontents: Essays on the History of Inductive Probability*. Cambridge Studies in Probability, Induction, and Decision Theory. Cambridge University Press, Cambridge, 2005.

[11] A. Orlitsky and N. Santhanam. Lecture notes on universal compression. Available online from `http://www-ee.eng.hawaii.edu/~prasadsn/`.